**A rational model of people's inferences about
others' preferences based on response times**

Vael Gates

Department of Psychology
University of California, Berkeley
2121 Berkeley Way, Berkeley, CA 94704

Frederick Callaway

Department of Psychology
Princeton University
South Dr, Princeton, NJ 08540

Mark K. Ho

Department of Psychology
Princeton University
South Dr, Princeton, NJ 08540

Thomas L. Griffiths

Department of Psychology
Princeton University
South Dr, Princeton, NJ 08540

Author Note

Corresponding author: Vael Gates, *vgates@berkeley.edu*.

Declaration of interest: None. Materials and data on Open Science Framework:

<https://osf.io/pczb3>

Abstract

There's a difference between someone instantaneously saying "Yes!" when you ask them on a date compared to "...yes." Psychologists and economists have long studied how people can infer preferences from others' choices. However, these models have tended to focus on what people choose and not how long it takes them to make a choice. We present a rational model for inferring preferences from response times, using a Drift Diffusion Model to characterize how preferences influence response time and Bayesian inference to invert this relationship. We test our model's predictions for three experimental questions. Matching model predictions, participants inferred that a decision-maker preferred a chosen item more if the decision-maker spent longer deliberating (Experiment 1), participants predicted a decision-maker's choice in a novel comparison based on inferring the decision-maker's relative preferences from previous response times and choices (Experiment 2), and participants could incorporate information about a decision-maker's mental state of cautious or careless (Experiments 3, 4A, and 4B).

Keywords: social cognition, inference, theory of mind, drift diffusion model

A rational model of people’s inferences about others’ preferences based on response times

It’s a quiet afternoon, and you and a new friend are trying to entertain yourselves. “Hiking? Movie?” you suggest. “Movie,” she says immediately. Feeling imaginative, you offer another choice: “Movie or science museum?” She takes longer to think on this. “...movie,” she says eventually. Taking into account how long it took your friend to make her choices, you might guess that she hates hiking but science museums are a good future activity.

When we observe others making choices, we learn from the concrete choices they make, but also from cues like body language and timing. Inferring preferences from response times has a long history in economics and psychology (e.g. Busemeyer, 1985; Busemeyer & Townsend, 1993; Chabris et al., 2008; Chabris et al., 2009; Diederich, 1997, 2003; Gill & Prowse, 2017; Moffatt, 2005; Spiliopoulos & Ortmann, 2018; Wilcox, 1993). Recent sequential sampling models solve this problem by thinking of response time as an observable measure of a noisy cognitive process, where evidence is accumulated over time and a choice is made between options when the integrated evidence crosses a threshold. The longer it takes to make a decision, the closer those items are in the decision-maker’s preferences, and the harder it is for the decision-maker to differentiate them. This principle has driven the use of sequential sampling models, the most common of which is the Drift Diffusion Model (DDM), to infer decision-makers’ preferences from their response times (Bogacz et al., 2006; Busemeyer & Rieskamp, 2014; Ratcliff, 1978; Ratcliff & McKoon, 2008; Ratcliff et al., 2016). Previous work has established that models such as the DDM can be used to infer people’s preferences from their response times (see section “Background and previous work” below).

Here we ask a different question: Can people infer *other people’s* preferences based on observing their response times and choices, modeling those others’ decision processes as a diffusion process? In modeling terms, the DDM predicts a decision-maker’s response times and choices given their preferences. Can we capture observers’ inferences using a rational model in which we invert the DDM to predict decision-makers’ preferences (Figure 1)?

A. Modeling RT based on preferences with the DDM **B. Inferring preferences from RT by inverting the DDM**

Figure 1. Schematic describing the Drift Diffusion Model (DDM) and the inverted DDM used in this paper. (A) The DDM is a generative model which predicts people’s response times (RT) and choices given their preferences across options. Parameters include the drift rate $\mu = \beta(u_a - u_b)$ (value difference, or strength of preference; u_a and u_b are the utilities for the two items), threshold θ (carefulness), and the drift multiplier β . The stronger the preference for an item, the faster the decision. (B) In this work, we invert the DDM to generate predictions for the inference task: people inferring others’ preferences based on observing their response times and choices.

People make inferences about preferences from response times constantly; using your friend’s response times and choices to infer her preference for science museums over hiking is one illustrative example. People can make even more complex inferences from response times and choices if they know the decision-maker’s mental state, like how carefully they are

making a decision. Your grandfather knows that if someone’s making a choice in a hurry, then their response times are less informative. He likewise knows that if someone’s very cautious, then a fast decision is particularly meaningful. In this paper, we propose that these forms of everyday inferences can be formalized and described by inverting a generative process like the DDM.

We explore the predictions of our inverse DDM through three experimental questions. In Experiment 1, we confirm the basic prediction that people will infer how much a decision-maker values one item over another based on observing the decision-maker’s response times. In Experiment 2, we extend this analysis to testing people’s predictions for a decision-maker’s novel choice based on having previously inferred their relative strength of preferences. Just as you could infer a new friend likes science museums more than hiking based only on two indirect questions, our model suggests that participants will predict a decision-maker’s choice in a novel comparison based on inferring the decision-maker’s relative preferences from previous response times and choices. In Experiments 3, 4A, and 4B, we explore the possibility that participants will be able to infer even more from a decision-maker’s response time if they are aware of the state of the decision-maker. When someone makes a quick decision, maybe it is because they value one choice much more than another, but maybe it is because they are feeling careless that day and are rushing through inessential decisions. These two factors correspond to different parameters of the DDM—“value difference” and “threshold”—which jointly predict the response time, and we ask if participants can perform joint inference over these parameters. We thus inform participants of the decision-maker’s carefulness, and observe whether participants’ predictions about their preferences vary.

Background and previous work

DDMs are often applied to perceptual decision-making tasks (e.g. Bitzer et al., 2014; Brunton et al., 2013). However, DDMs have also been used to model value-based decisions

integrating both choice and timing information (Shadlen and Shohamy (2016) and Polania et al. (2014) discuss relationships between DDMs describing perceptual versus value decision tasks). An important property of DDMs applied to value-based decisions is that they predict that response time should correspond to strength of preference, meaning that how much a decision-maker values one item over another (“value difference”) should be expressed in response time, where smaller value differences correspond to longer response times (Echenique & Saito, 2017). This prediction has been supported in a variety of settings: risky choice (Alós-Ferrer & Garagnani, 2020; Konovalov & Krajbich, 2019), intertemporal choice (Amasino et al., 2019; Dai & Busemeyer, 2014; Konovalov & Krajbich, 2019; Krajbich, Bartling, et al., 2015), social decision-making (Frydman & Krajbich, 2018; Konovalov & Krajbich, 2019; Krajbich, Hare, et al., 2015; Krajbich et al., 2014), and food choice (Clithero, 2018a, 2018b; Krajbich et al., 2010; Krajbich & Rangel, 2011; Milosavljevic et al., 2010; Towal et al., 2013).¹

Previous work has demonstrated that people can empirically infer the preferences of others from their response times. The DDM specifically has been used as motivation for some of these studies, though the authors do not formally invert the DDM as we do. For example, in Konovalov and Krajbich (2020), participants in a strategic bargaining experiment inferred their opponents’ preferences based on both explicitly stated and observed response times, and changed their behavior to make use of this knowledge. In another study, participants in Frydman and Krajbich (2018) played a strategic information cascade game, in which they had to determine the binary state of the environment, receiving

¹ Previous work has studied response times in the framework of dual-process theory: If decisions are fast, this suggests that an intuitive process system is being used, and slow decisions suggest that a slow, deliberative process is being used (Rand et al., 2012; Rubinstein, 2007). Here, we interpret response time differences as representing differences in strength of preferences, on the basis of other work that shown that findings interpreted through the dual-process theory lens are consistent with single-process strength-of-preference predictions (Konovalov & Krajbich, 2020; Krajbich, Bartling, et al., 2015; Krajbich, Hare, et al., 2015; Zhao et al., 2019).

both a probabilistic private signal about the environment and also observing other players' public choices. The DDM predicts that choices should be slow when players' private signals were in conflict with other players' public choices, which the authors observed. The study then examined how people behaved when presented with other players' response times. Participants used this information to infer how unsure other players were about their decisions: When public choices were incorrect, participants in the response-time condition were significantly more likely to follow their private signals compared to participants without access to other players' response times. Finally, an interesting consequence of the DDM's prediction that people will spend more time on harder decisions is that participants can act suboptimally under circumstances with a fixed time limit and payoffs allocated per-choice. Participants may spend a suboptimal amount of time on low-stakes choices with small value differences rather than high-stakes choices with large value differences. When researchers observed this result, they found that by forcing a cutoff time for all decisions they could significantly improve participants' payoffs (Krajbich et al., 2014; Oud et al., 2016).

Unlike previous work, we invert the DDM to predict how people will infer preferences from other decision-makers' choices and response times, fitting our model with participants' preference estimates about the decision-maker. This use of Bayes' rule to invert a decision-making model follows Lucas et al. (2014), Jara-Ettinger et al. (2016), Baker et al. (2017), and Jern et al. (2017), who inverted choice models to describe how people infer preferences from the outcomes others select. We continue to follow their tradition of using inverse decision-making models to explore theory of mind, extending this work to incorporate response times.

Bayesian inference over preferences from response times

We begin with a broad overview of our inverted Drift Diffusion Model (DDM), in which we use Bayesian inference to infer preferences from response times. The basic setup is that one person (the *observer*) watches a second person (the *decision-maker*) make a choice

between two options. The observer then infers a distribution over how much the decision-maker likes each option (their *inferred utilities*) based on the option they chose and—critically—the amount of time they took to make the decision.

To make such an inference, the observer must have a *generative model* of the decision-maker’s decision-making process. Here, we assume that this model is a DDM (Figure 1).² In this model, decisions are made on the basis of sequentially accumulated noisy evidence about the difference of the two options’ utilities. The accumulation process runs until the total evidence in favor of one option or the other exceeds a threshold, at which point, the corresponding option is chosen. Formally, the evidence, x , is a dynamical system with dynamics

$$dx = \mu + \sigma dW, \quad (1)$$

where μ is the *drift rate*, σ specifies the amount of noise in the integration, and W is a Wiener process (the continuous limit of a Gaussian random walk, also called Brownian motion). The evidence is initialized at zero. Following Milosavljevic et al. (2010), we assume that the drift rate is a linear function of the difference in values of the two items,

$$\mu = \beta(u_a - u_b), \quad (2)$$

where u_a and u_b are the utilities of the two items being chosen between, and the drift multiplier, β , can be interpreted as the decision-maker’s sensitivity to differences in value. The accumulation process continues until the evidence crosses one of the decision boundaries. We assume symmetric constant boundaries, at a distance of θ from the initial point of zero. Intuitively, θ controls how careful the decision-maker is: larger values make decision slower but more accurate. If at any moment $x > \theta$, option a is chosen; we denote this event $a \succ b$. If $x < -\theta$, option b is chosen, that is, $b \succ a$. The time point at which this event occurs, t , is

² We chose this model because it is commonly used in decision-making research and is computationally easy to work with. However, the framework can accommodate any model that defines a joint distribution over choices and response times.

the response time (for simplicity, we do not consider the non-decision time variable that is often added to t to produce the response time). The DDM thus defines a probability distribution

$$p_{\text{DDM}}(a \succ b, t \mid u_a - u_b; \beta, \theta). \quad (3)$$

To infer the decision-maker’s preferences given the observed choice and response time, the observer must invert this model. We can do this using Bayes rule, resulting in a posterior distribution over the utilities

$$p(u_a, u_b \mid a \succ b, t) \propto p(u_a)p(u_b) \cdot p_{\text{DDM}}(a \succ b, t \mid u_a - u_b; \beta, \theta), \quad (4)$$

where $p(u_a)$ and $p(u_b)$ capture the observer’s prior distribution over utilities. For simplicity, we assume a standard Gaussian prior.

Equation 4 specifies how a rational agent should update their beliefs about another person’s preferences based on an observed choice and response time. However, the inferences one draws depend on the parameters of the DDM: the drift multiplier, β , and the threshold, θ . Intuitively, these parameters can be thought of as individual difference or mental state variables that jointly determine the accuracy and speed of the decision. We explore the θ variable further in Experiment 3. For now, it is sufficient to note that the qualitative model predictions in Experiments 1 and 3 are insensitive to these parameters. However, for the purpose of plotting the predictions, we fit the parameters by minimizing the sum of squared errors between the model predictions and aggregate participant responses across the first two experiments. We treat the third experiment separately, as described in the methods for that experiment.

To summarize, the DDM serves as a generative model that relates decision-making parameters (the drift rate and threshold) to choices and response times. Critically, an observer can then invert this model to make inferences about the decision-maker’s preferences. Here, we propose that human inferences from choices and response times will be consistent with such an ideal Bayesian observer model. To illustrate in a concrete example,

recall our introductory example in which your friend takes longer to choose “movie” when it is paired with science museum than when it is paired with hiking. In considering her responses, you would conclude not only that she really wants to see a movie, but that she seems to like science museums more than hiking. Remarkably, this intuitive inference falls naturally out of inverting the DDM: Her longer response time in the movie-museum pairing reflects a more shallow drift rate, which results from the movie-museum utility difference being smaller than the movie-hiking difference. Since the movie has the highest utility (after all, she chose it twice), this means that the museum has higher utility than hiking.

In the remainder of the paper, we empirically test three predictions of our model. Experiment 1 tests the prediction that observing faster choices indicates that the decision-maker has a larger preference difference. Our inverse DDM predicts this because a large difference in utilities between two items will result in a steeper drift rate, which leads to a faster decision. In Experiment 2, we test the prediction that people can draw inferences about novel choice pairs even when the choices alone provide no information about the relative utility. In particular, we present participants with scenarios similar to the movie/hiking/museum example and test whether they integrate information across observations consistent with inverting the DDM. Finally, in Experiments 3, 4A, and 4B, we examine whether people’s conclusions about preferences from response times is affected by background knowledge about the decision-maker’s mental state, specifically whether the decision-maker is feeling cautious or careless. This corresponds to performing inference about the decision-maker’s drift rate (i.e., the relative utilities of the options) conditioned on different decision thresholds.

Experiment 1: Inferring preferences from response times

We test the core prediction of the DDM that decisions will be faster when the difference in preference for the two items is larger. In the model, a larger preference difference results in a higher drift rate and more rapid movement towards a decision

boundary. Thus, if people employ a model like the DDM when inferring other’s preferences from their response times, they should rate a chosen item as being more strongly preferred when the decision is made more quickly.

Methods

Participants. 481 participants with U.S. IP addresses were recruited through Amazon Mechanical Turk. Participants were paid \$1.00 in compensation. Participants were excluded from the study if, for the critical questions, they indicated the decision-maker preferred a non-chosen item more than once, indicating lack of attention. 66 participants were excluded for a total of 415 participants.

We used power analyses to determine the sample size for Experiments 1, 2, and 3 based on effect sizes observed in pilot studies. These power analyses, our exclusion criteria, our stimuli, and our planned procedures and analyses were preregistered here: <https://osf.io/8n3kd>. Participants could only complete one of the pilot experiments or main experiments.

Experiments 1, 2, and 3 were IRB-approved by University of California, Berkeley, Committee for Protection of Human Subjects / Office for Protection of Human Subjects, Protocol ID: 2015-05-7551, Protocol Title: Cognitive Research Using Amazon Mechanical Turk (Expedited), with all participants giving informed consent.

Stimuli. In all experiments, participants viewed surveys created on Qualtrics containing text and videos. Experiments 1, 2, and 3 used the same set of videos (or a subset of these videos), which were created by filming a decision-maker making eight choices. The videos began by the decision-maker pulling two pieces of paper apart, which were covering two labeled bowls (e.g. A and B; bowls could be labeled A, B, or C). The decision-maker thought about their options while staring between the two bowls, and then they reached into one of the bowls to make a choice. The decision-maker made choices after 3, 5, 7, and 9 seconds after video onset, and made choices with either their left or right hand, for a total of

eight videos. To counterbalance asymmetries in the videos, the number of videos used was doubled by creating flipped copies over the vertical axis, for a total of 16 videos. To counterbalance the effects of A, B, and C labels, the labels of the 16 videos were edited to use the label combinations AB, BA, CA, AC, BC, and CB, for a total of 96 videos. The text was modified to match the video labels. For narrative simplicity, we will only refer to the AB label combinations. Stimuli for all experiments are available on OSF: <https://osf.io/pczb3>.

Procedure. After consenting, participants were shown introductory text: “In this experiment, you will see videos of someone choosing between an item in one bowl and an item in another bowl. There are 8 videos. Please do your best to answer the questions afterward.”

Participants then read: “Two items, A and B, are inside the bowls below. This person has just been asked to choose which item they want. Please watch the video.” Participants then watched a video in which a decision-maker chose an item from bowl A or B within 3, 5, 7, or 9 seconds from the onset of the video. After watching the video, participants were asked: “What are the person’s feelings about item A and item B?” Participants then moved a slider containing values from 0 (labeled “Strongly Prefers A”) to 100 (labeled “Strongly Prefers B”) and set initially at 50 (labeled “Neutral Between A and B”). No grid lines were shown, nor the numbers 0-100, and participants were required to click and move the slider (but they could move it back to 50 if desired). After answering this question the participant could advance to the next trial, which displayed on another page (Figure 2). Participants could replay videos as often as they wished, and all of the text/video/questions were on the same page. Participants could not return to a previous page. There were eight trials: Participants saw a total of eight videos (A/B choice $\times$ 3/5/7/9 seconds) in random order.

We will refer to the choices as between A and B for narrative simplicity. However, participants were randomly assigned to one of 12 counter-balancing variants in which the labels and video orientations were randomized. Within every variant, the labels were held constant (options: AB, BA, AC, CA, BC, CB) and the video orientation was held constant (the videos were either as filmed, or were flipped over the vertical axis).

Two items, A and B, are inside the bowls below. This person has just been asked to choose which item they want.

Please watch the video.

What are the person's feelings about item A and item B?

Strongly Prefers A

Neutral Between A and B

Strongly Prefers B

>>

Figure 2. Experiment 1 Stimulus. Participants watched a 3, 5, 7, or 9 second video of the participants choosing between A or B, then answered the question below.

After watching all eight videos and answering the accompanying questions, participants had the opportunity to write comments then exited the survey.

Model. We assume that participant responses are based on the posterior mean estimate of the difference in utilities of the two items. This is defined as

$$\mathbb{E}[u_a - u_b \mid a \succ b, t] = \int \int (u_a - u_b) \cdot p(u_a, u_b \mid a \succ b, t) du_a du_b. \quad (5)$$

However, because the DDM likelihood depends only on the difference in values (Equation 3), we can re-express Equation 5 as a single integral over the difference,

$$\mathbb{E}[\delta_u \mid a \succ b, t] = \int \delta_u p(\delta_u \mid a \succ b, t) d\delta_u, \quad (6)$$

where $\delta_u = u_a - u_b$ and

$$p(\delta_u \mid a \succ b, t) = \frac{1}{Z} \text{Normal}(\delta_u; \mu = 0, \sigma^2 = 2) \cdot p_{\text{DDM}}(a \succ b, t \mid \delta_u). \quad (7)$$

Note that the variance of the difference between two Gaussians is the sum of the variance of each, hence $\sigma^2 = 2$. We have suppressed β and θ for concision. We compute the normalizing constant, Z , and the integral in Equation 6 by adaptive numerical integration (Genz & Malik, 1980).

To convert this utility difference (which is in arbitrary units) to the scale defined by the slider participants used to make a response, we use a scaling parameter, such that the predicted response is $\alpha \mathbb{E}[u_a - u_b \mid a \succ b, t]$. We fit α (along with β and θ) to aggregate participant behavior by minimizing the sum of squared errors.

As one would expect, the DDM predicts that a choice of item a over item b is increasingly likely as $u_a - u_b$ increases. However, the choice alone only weakly constrains the size of the difference. As illustrated in Figure 1, the response time, t , provides much more information. With a strong preference, a fast response is quite likely and a slow response is quite unlikely. With a weak preference, probability is more evenly spread across response times. As a result, when we invert the DDM, we find that a strong preference is more likely given a fast response (because a fast response is likely given a strong preference) and a weak preference is more likely given a slow response (because a slow response is unlikely given a strong preference).

Results

In Experiment 1, we showed participants videos in which a decision-maker took 3, 5, 7, or 9 seconds to decide between two items, then asked participants about the decision-maker's preferences about the items. We expected that participants would infer that the decision-maker preferred the chosen item more when the response time was shorter (e.g. 3 seconds) compared to when the response time was longer (e.g. 9 seconds), as predicted by

our model. This is indeed what we observed, with a qualitatively close match between the model’s predictions and the results (Figure 3).

Figure 3. Experiment 1 Results. (a) Experimental Results. The mean $\pm$ SE of inferred preference for the chosen item, $n=415$, is shown for each of the response time conditions. Responses were reoriented to be 0-50: 0 indicates the participant felt the decision-maker was neutral between items A and B, and 50 indicates the participant felt the decision-maker strongly preferred the chosen item (A or B, whichever was appropriate to the trial). Light grey lines indicate individual participants’ averaged values for each time point. (b) Model Predictions.

To analyze the effect of response time, we created a linear mixed effects model with response time as a continuous fixed effect, subject as a random effect, and participants’ preference judgments (0-100, 50 was “Neutral Between A and B,” 0 was “Strongly Prefers A” and 100 was “Strongly Prefers B”) as the continuous output variable. Each participant had eight data points: two answers each for the 3, 5, 7, and 9-second questions. The estimate (beta parameter) for the fixed effect of response time was -2.87 (95%CI=[-3.03,-2.71], $t(2904.0)=-35.16$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.30$). (The estimate for the fixed effect intercept was 44.91, 95%CI=[43.34,46.48], $t(1003.4)=56.10$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.76$.) The estimated power of the response time predictor

for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R package “simr.”

Discussion

As predicted by our model, participants inferred that a decision-maker preferred an item more the less time they spent selecting the item. This effect was statistically significant, and qualitatively matched our model results. Interestingly, the individual participant responses in Figure 3 suggest that the main effect of response time may be even stronger than indicated by the mean responses, because a minority of participants ($40/415 = 9.6\%$) do not take the decision-maker’s response time into account. Instead, they inferred that the decision-maker had the maximal preference for the chosen item on every question (shown in Figure 3 by a set of $\{50, 50, 50, 50\}$ responses, indicating the participant moved the slider to 100 on four trials and 0 on four trials).

Experiment 2: Predicting novel choices based on relative strength of preferences

In Experiment 1, we demonstrated that participants could infer preferences of a decision-maker from their choices and their response times. However, people regularly perform more complex inferences than that: If your friend did not hesitate at all before choosing movies over hiking, but hesitated a while before choosing movies over science museums, then you could guess that she preferred science museums to hiking. We now explore this ability to predict a decision-maker’s choice in a novel comparison, based on having inferred their relative strength of preferences from response times in previous choices.

Inferring graded preferences from choices (or response times) that will allow prediction for novel comparisons is not often studied. The most closely related literature surrounds *transitive inference*. In transitive inference paradigms, participants—children, adults or animals—are asked to draw conclusions like if $B > A$, and $A > C$, then $B > C$. In these paradigms, associations are often manually taught between previously-unassociated items or

assertions, and transitive inference is solely a function of choices (e.g. Acuna et al., 2002; Harris & McGonigle, 1994; Maybery et al., 1986; Thayer & Collyer, 1978). It is rare to show participants executing transitive inference about preferences, which are distinct in that they are often socially learned through implicit cues in addition to appearing probabilistic and non-causal. Hu et al. (2015) is an exception, demonstrating that young children can perform indirect, graded transitive inference of another decision-maker’s preferences based on observing their choices. Children in this study inferred that a puppet preferred object B over C ($B > C$), after watching the puppet choose B consistently over A ($B \gg A$) and the puppet choosing C somewhat consistently over A ($C > A$).

Similar to Hu et al. (2015), we test whether people can infer another decision-maker’s graded preferences by observing their choices. However, rather than watching the decision-maker make choices repeatedly, we had participants watch the decision-maker make two decisions. We investigate whether people can infer a decision-maker’s preferences enough to predict their selection on a novel choice after only having observed the decision-maker’s choices and response times twice. We can do this because it can be much more efficient to use response time and choice information together than choice information alone. If A were chosen over B, and A were chosen over C, people would have to observe more choices to statistically infer the graded preferences that would allow them to make a prediction for an unseen choice between B and C. Using response time information, people could make this prediction immediately, given they had inferred the decision-maker’s relative strength of preferences after observing their previous choices and response times.

In Experiment 2, participants saw pairs of videos. In the first video, the decision-maker made a choice between two items, A and B, within 3 seconds. In the second video, the decision-maker made a choice between the original item and another item, A and C, within 9 seconds. Participants were then told that the decision-maker had to make a choice between B and C, and were asked what choice they thought the decision-maker would make and how likely that choice was. Our model predicts that participants would anticipate the

decision-maker’s selection on the novel choice, having inferred the decision-maker’s preferences based on response times and choices from the previous two decisions. Our model also proposes that participants’ inferred likelihoods of the predicted choice would be lower than in the control condition in which participants only needed to use choices to make the prediction (and did not need to take response time into account).

Methods

Participants. 478 participants with U.S. IP addresses were recruited through Amazon Mechanical Turk. Participants were paid \$1.50 in compensation. No exclusion criteria were applied. The sample size was determined by a power analysis based on the effect size of pilot studies and was preregistered.

Stimuli and Procedure. Video stimuli were the same as in Experiment 1.

After consenting, participants read introductory text: “In this experiment, you will see videos of someone choosing between an item in one bowl and an item in another bowl. There are 8 sets of videos. Please do your best to answer the questions afterward.” Participants then entered the main experiment.

In Experiment 2, participants saw pairs of videos. Participants saw the text “This person is choosing between two items, A and B,” and then were shown a first video in which the decision-maker makes a choice within 3 seconds. On the same page, participants were then presented with the text “Now they are choosing between A and C,” and shown another video of a decision-maker making a choice within 9 seconds. The participant was then shown, also on the same page, “This person is now offered a choice between items B and C. What choice do you think they’d make, and how likely do you think it is that they’d make that choice?” Participants then moved a slider containing values from 0 (labeled “Very likely B”) to 100 (labeled “Very likely C”) and set initially at 50 (“Equally likely”). No grid lines were shown, nor the numbers 0-100, and participants were required to click and move the slider (but they could move it back to 50 if desired). After watching the two videos and answering

This person is choosing between two items, A and B.

Now they are choosing between items A and C.

**This person is now offered a choice between items B and C.
What choice do you think they'd make, and how likely do you think it is that they'd make that choice?**

Very likely B

Equally likely

Very likely C

>>

Figure 4. Experiment 2 Stimulus. Participants watched a 3-second video of the decision-maker deciding between A and B, and a 9-second video of the decision-maker deciding between A and C. Participants then indicated whether they believed the decision-maker would be most likely to choose B or C in a novel choice.

the question, participants could advance to the next page, where they watched another pair of videos and answered another question (Figure 4). There were eight pages/questions total. Participants could replay videos as often as they wished, and all of the text/videos/questions were on the same page. Participants could not return to a previous page. After watching all eight pairs of videos and answering the accompanying questions, participants had the opportunity to write comments then exited the survey.

For narrative simplicity, we will refer to the first video as a 3-second choice between A and B, and the second video as a 9-second choice between A and C. However, randomized counter-balancing variants were used in the experiment. We randomized the order of which video was presented first within-participant (half of the trials showed the 3-second video first, and half of the trials showed the 9-second video first). Between-participants, we randomized the order of the final choice between participants (in one variant, participants made a choice between B and C, and in another variant participants made a choice between C and B), the labels (one label was kept consistent on one side, resulting in the following variants: AB vs AC, BA vs CA, BA vs BC, AB vs CB, CA vs CB, and AC vs BC), and video orientation (half of the variants had all videos flipped over the y-axis). This resulted in a total of 24 variants ($2 \text{ final choice options} \times 6 \text{ label pairs} \times 2 \text{ video orientations}$), and 8 questions per participant ($2 \text{ orderings of 3/9-second videos} \times 4 \text{ choice options (R/R, R/L, L/R, R/R, where "R/L" represents the decision-maker choosing the right-most object in the first video and the left-most object in the second video)}$)).

Participants were randomly assigned to one of these 24 variants, within which they viewed the 8 pairs of videos wherein labels and video orientation were held constant but response time and choices varied. A given participant would be shown the video pairs in random order, but the videos within the pairs were ordered.

Analysis. Our stimuli resulted in four conditions. Participants watched the decision-maker choose between A and B in one choice, then A and C in another, and in the critical question participants were asked whether the decision-maker would prefer B or C in

a novel choice. We thus defined the four conditions based on whether the decision-maker chose B or C in either of the observed choices. The condition “NeitherChosen” indicated that the decision-maker never chose B or C in the observed choices (choosing A instead), and the condition “BothChosen” indicated that the decision-maker chose B in one choice and C in another. The condition “FastChoice” indicated that the decision-maker chose B once, and “Slow choice” indicated that the decision-maker chose C once (Figure 5). Each participant answered two questions for each of the four conditions, for a total of eight data points.

Choice, 3sec	A > B	B > A	B > A	A > B
Choice, 9sec	A > C	C > A	A > C	C > A
Predicted Choice	<u>C > B</u>	<u>B > C</u>	<u>B > C</u>	<u>C > B</u>
Condition	NeitherChosen	BothChosen	FastChoice	SlowChoice

Figure 5. Experiment 2 Conditions. Participants watched a 3-second video of the decision-maker deciding between A and B, and a 9-second video of the decision-maker deciding between A and C. In the critical question, participants indicated whether they believed the decision-maker would be more likely to choose B or C in a novel choice. The model’s predictions for participants’ predicted choices are shown. Four conditions were defined based on whether the decision-maker chose B or C in either of their choices (3-second and 9-second). In the “NeitherChosen” condition, the decision-maker chose neither B nor C, and in the “BothChosen” condition, the decision-maker chose B and C. For the non-inference control conditions, in the “FastChoice” condition, the decision-maker chose B quickly, and in the “SlowChoice” condition, the decision-maker chose C slowly.

The first two conditions, “NeitherChosen” and “BothChosen,” require that the participant make an inference based on response time, since in the “NeitherChosen” condition the decision-maker never preferred B or C in the observed choices, and in the “BothChosen” condition the decision-maker preferred both. Thus when the participant was asked whether the decision-maker preferred B or C, the answer was ambiguous based solely on the choices the participant had seen.

The second two conditions, “FastChoice” and “SlowChoice,” did not require that the participant make an inference based on response times, just choices, since the decision-maker preferred either B or C in the choices they made in the videos. Thus when the participant was asked whether the decision-maker preferred B or C, the decision-maker should logically choose whichever choice the decision-maker preferred in one of the videos.

Model. We assume that participants respond with the posterior mean probability of the unseen choice given the two observed choices and response times. To compute this value, we first compute a posterior over each item’s value, and then integrate over these values to produce a predicted choice. Here, we specify these steps for the “NeitherChosen” condition; the derivations for the other conditions have the same form. Let $\mathcal{D} = \{a \succ b, t_{ab}, a \succ c, t_{ac}\}$ denote the observed data. The posterior over values is then given by

$$p(u_a, u_b, u_c \mid \mathcal{D}) \propto \varphi(u_a)\varphi(u_b)\varphi(u_c) \cdot p_{\text{DDM}}(a \succ b, t_{ab} \mid u_a - u_b) \cdot p_{\text{DDM}}(a \succ c, t_{ac} \mid u_a - u_c), \quad (8)$$

where φ is the standard normal pdf, and the posterior over choice is given by marginalizing over the three values as well as the predicted decision time,

$$p(b \succ c \mid \mathcal{D}) = \int \int \int \int_0^\infty p_{\text{DDM}}(b \succ c, t_{bc} \mid u_b - u_c) \cdot p(u_a, u_b, u_c \mid \mathcal{D}) dt_{bc} da db dc. \quad (9)$$

At a high level, the key model prediction is that people will be able to predict a novel choice even when the previous choices (without response times) do not provide any information about which option is preferred. Consider the “NeitherChosen” case. Here, the observer must predict a choice between B and C after observing A being chosen over both of them. This implies that $u_a > u_b$ and $u_a > u_c$, but provides no information about $u_b - u_c$. However, considering the response times of each choice (short and long, respectively) provides additional information about the relative strength of preference in the first two choices. Following the predictions of Experiment 1, the observer infers that $u_a - u_b$ is large because this decision was made quickly and that $u_a - u_c$ is small because this decision was made slowly. Thus, u_c must be greater than u_b , and C is more likely to be chosen over B.

This line of reasoning can be summarized in notation as:

$$a \succ b \wedge a \succ c \wedge t_{ac} > t_{ab} \implies u_a - u_c < u_a - u_b \implies u_c > u_b \implies p(c \succ b) > \frac{1}{2}. \quad (10)$$

The logic is similar in the “BothChosen” case, but inverted:

$$b \succ a \wedge c \succ a \wedge t_{ac} > t_{ab} \implies u_c - u_a < u_b - u_a \implies u_c < u_b \implies p(b \succ c) > \frac{1}{2}. \quad (11)$$

For comparison, we also consider cases where the previous choices alone *do* constrain the prediction about the novel choice, by transitivity of preference. In the “FastChoice” condition we have

$$b \succ a \wedge a \succ c \implies u_b > u_a \wedge u_a > u_c \implies u_b > u_c \implies p(b \succ c) > \frac{1}{2}, \quad (12)$$

and in the “SlowChoice” condition,

$$a \succ b \wedge c \succ a \implies u_a > u_b \wedge u_c > u_a \implies u_c > u_b \implies p(c \succ b) > \frac{1}{2}. \quad (13)$$

The model makes a stronger prediction in the “FastChoice” and “SlowChoice” conditions because choices provide stronger and more direct evidence of preference ordering compared to differences in response time.

Results

Our model predicts that participants would anticipate a decision-maker’s novel choice based on inferring the relative strength of that decision-maker’s preferences from previous response times and choices. Specifically, participants should predict the decision-maker’s choice better than chance. Moreover, participants’ inferred likelihoods of the predicted choice should be lower than the control conditions in which only choices (and not response times) were necessary. We compare our results qualitatively to these model predictions (Figure 6).

Difference from chance, all conditions. To determine whether participants could make predictions for a novel choice better than chance, we conducted the equivalent of

Figure 6. Experiment 2 Results. (a) Experimental Results, Mean $\pm$ SE of Inferred Likelihood Estimates for the Predicted Choice, $n=478$. Shown in blue are participants’ averaged responses for each of the conditions; the line at 50 indicates equal likelihood between choices (chance); black points indicate individual participant averaged responses for each condition. Participants saw the decision-maker choose between A and B, then A and C, and had to infer whether the decision-maker preferred B or C in the the unseen choice (the model’s predicted choices are shown). The “NeitherChosen” and “BothChosen” conditions required that participants infer the decision-maker’s likelihood of preferring an item from choices and response time, while the “FastChoice” and “SlowChoice” conditions only required that participants infer from the decision-maker’s choices. (b) Model Predictions.

one-sample t -tests for a linear mixed effects model.³ Condition was a categorical fixed effect (“NeitherChosen”/“BothChosen”/“FastChoice”/“SlowChoice”), subject was a random effect, and likelihood judgment was a continuous output variable.

In the “NeitherChosen” condition, we predicted better-than-chance performance: specifically that if the decision-maker chose $A \succ B$ quickly, and $A \succ C$ more slowly, that

³ Specifically, we translated participants’ likelihood judgments to be between -50 and 50 (where 0 represented a judgement that it was equally likely the decision-maker would choose either of the two items) then did not fit the fixed effect intercept (which defaults to 0).

participants would infer that the decision-maker had preferences $A > C > B$. From our linear mixed effects model described above, the estimate (beta parameter) for the “NeitherChosen” condition was 7.99 (95%CI=[5.89,10.10], $t(1681.5)=7.46$, $p=1.4e-13$, t -tests using Satterthwaite’s method, $\eta_p^2=0.03$),⁴ in the expected direction $C > B$. The estimated power of the “NeitherChosen” condition against chance for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a t -test, generated via 200 simulations with the R package “simr.”

In the “BothChosen” condition, we expected that if the decision-maker chose $B \succ A$ quickly, and $C \succ A$ slowly, then participants would infer that the decision-maker had preferences $B > C > A$ above chance. From our linear mixed effects model described above, the estimate (beta parameter) for the “BothChosen” condition was 14.53 (95%CI=[12.43,16.63], $t(1681.5)=13.55$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.10$), in the expected direction $B > C$. The estimated power of the “BothChosen” condition against chance for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a t -test, generated via 200 simulations with the R package “simr.” These results match our model predictions, with one difference being that the model predicted a more similar level of inference for the “NeitherChosen” and “BothChosen” conditions than was observed empirically.

We also expected that participants would make the appropriate choices in the “FastChoice” and “SlowChoice” control conditions, that they would do so above chance performance, and that they would do so more accurately than in the “NeitherChosen” and “BothChosen” inference conditions. The “FastChoice” and “SlowChoice” conditions act as control conditions: In these conditions, participants only needed to use choice information to correctly infer the decision-maker’s preferences, instead of also needing response times. From our linear mixed effects model, the estimate (beta parameter) for the “FastChoice” condition was 25.17 (95%CI=[23.07,27.27], $t(1681.5)=23.47$, $p<2e-16$, t -tests using Satterthwaite’s

⁴ All η_p^2 effect sizes were calculated using the assumption that $t^2 = F$, followed by $\eta_p^2 = F * df_{num} / (F * df_{num} + df_{res})$ (e.g. Lakens, 2013) with $df_{num} = 1$.

method, $\eta_p^2=0.25$), in the expected direction $B > C$. The estimated power of the “FastChoice” condition against chance for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a t -test, generated via 200 simulations with the R package “simr.” The estimate for the “SlowChoice” condition was 22.34 (95%CI=[20.24,24.44], $t(1681.5)=20.84$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.21$), in the expected direction $C > B$. The estimated power for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a t -test. These results matched our model predictions. One difference was that the model predicted slightly more similar level of inference across these conditions than was observed empirically.

The “FastChoice” and “SlowChoice” results serve as an baseline for accuracy based purely on choice compared to the “NeitherChosen” and “BothChosen” results, which required inferences based both on choice and response time. We explore this comparison directly next.

Comparison between conditions: timing-inference conditions vs choice control conditions. In addition to examining how our results differ from chance, we also examine how the timing-based inference conditions (“NeitherChosen”/“BothChosen”) compare to the control conditions (“FastChoice”/“SlowChoice”). To this end, we constructed a linear mixed effect model where condition was a categorical fixed effect, subject was a random effect, and participants’ likelihood judgments were a continuous output variable. However, unlike the previous model, the conditions were merged (“NeitherChosen/“BothChosen” and “FastChoice”/“SlowChoice”) to create two comparison groups, and these groups were compared to each other.

The estimate (beta parameter) for the fixed effect of group (timing-inference conditions versus control conditions) was -12.49 (95%CI=[-14.22,-10.77], $t(3345.0)=-14.17$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.06$). (The estimate for the fixed effect intercept was 73.76, 95%CI=[72.04,75.47], $t(842.1)=84.24$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.89$.) The estimated power of the group predictor for $\alpha=0.01$ was 100.0% (95%CI=[98.17,100.0]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R package “simr.” Our empirical results matched model results,

showing a statistically significant pairing between the timing-based inference conditions and the control conditions, with the control conditions showing higher inferred likelihoods for predicted choice.

Discussion

In Experiment 2, we asked whether participants could make a prediction about an unseen choice based on inferring the relative strength of a decision-maker’s preferences from response time and choice information from two observed choices. Our model results and empirical results were in close alignment: In all conditions, participants inferred the predicted choice above chance. Moreover, participants had lower inferred likelihoods for the predicted choice for the conditions in which they needed to use both response time and choice information (the “NeitherChosen” and “BothChosen” inference conditions) compared to when they only needed to use choice information (the “FastChoice” and “SlowChoice” conditions, which serve as an baseline for accuracy based purely on choice).

Specifically, in the “NeitherChosen” condition, participants saw the choice $A \succ B$ made quickly, the choice $A \succ C$ made slowly, and the model predicted that they would infer that $C > B$, based on the DDM reasoning that items closer in value take a longer time to decide between. In the “BothChosen” condition, participants saw the choice $B \succ A$ made quickly, the choice $C \succ A$ made slowly, and the model predicted that they would infer that $B > C$, again based on the reasoning that items closer in value take a longer time to decide between. Interestingly, participants were better able to perform the inference for the “BothChosen” condition compared to the “NeitherChosen” condition, while the model predicted very similar inferred likelihoods for both of these “inference” conditions. Intuitively, it can feel easier to choose whether B or C was better if either of these items were chosen compared to if neither were chosen, but this intuition was not captured by the model. Future versions of the model should capture this more subtle behavior, perhaps by incorporating a bias for reduced timing for chosen items.

In the non-inference control conditions, participants inferred the predicted choice similarly across the two conditions “FastChoice” (the participants saw $B \succ A$ chosen quickly, $A \succ C$ chosen slowly, and had to infer $B \succ C$) and “SlowChoice” (the participants saw $A \succ B$ chosen quickly, $C \succ A$ chosen slowly, and had to infer $C \succ B$), as predicted by the model. Participants did have slightly higher inferred likelihoods for the predicted choice on the “FastChoice” condition, in which the target item was chosen quickly rather than slowly, which was not represented in the model. This was a relatively minor effect, but future work could to explore whether this element of human psychology could be captured in a modified DDM or with a different set of well-justified parameters. Relevantly, the model predictions for Experiment 2 were highly varied based on the parameter settings, so it could be interesting to explore what kind of behavior is described under those different settings.

Experiment 3: Sensitivity to a decision-maker’s mental state

We next investigated whether people would incorporate a decision-maker’s mental state when reasoning about response times and implied preferences. If a decision-maker quickly chose between two items, it could have been because one of the options was obviously preferred. However, if the decision-maker was feeling careless rather than cautious that day, their quick decision-making could also be attributed to being tired and rushing through decisions. Both factors—the value difference between the two choices, and the decision-maker’s overall carefulness—could influence response time. We asked whether participants would incorporate the decision-maker’s carefulness in reasoning about their preferences.

From a modeling perspective, in the DDM, a decision-maker’s carefulness is naturally modeled by the decision threshold θ , which determines the amount of evidence that must be accumulated in favor of an item before a decision is made. The higher the threshold, the less sensitive the decision-maker will be to random fluctuations in the evidence, and thus the less likely they will be to make an error by selecting the item that they actually disprefer.

However, this robustness comes at a cost: Accumulating evidence takes time, and so a decision-maker with a high threshold will make slower decisions. Importantly, the decision time depends on both the drift rate and the threshold. A slow decision can result from either a low drift rate (small preference difference) or a high threshold (high caution) (Figure 7).

The combined influence of drift rate and threshold on decision time results in a critical prediction of our model: knowing a decision-maker’s carefulness should affect one’s inferences about their preferences. Specifically, if you believe that a decision-maker is very cautious (has a high threshold), and you observe them make a fast decision between two items, you should infer a strong preference for the chosen item, because a high threshold is unlikely to be reached quickly unless the evidence is strong. In contrast, if you believe the decision-maker to be careless (low threshold), a fast decision is consistent with moderate or even no preference difference.

In Experiment 3, participants watched decision-makers making decisions quickly (3 seconds) or slowly (9 seconds), having been described as cautious (high threshold) or careless (low threshold). Participants were then asked how much they believed the decision-maker valued their chosen item. We hypothesized based on model predictions that if a decision-maker made a choice very quickly and was described as being cautious, then participants would infer that the decision-maker valued one of the items a lot more than the other. If a decision-maker made a choice very quickly but was described as careless, we expected that participants would attribute this quickness only partly to how much the decision-maker valued the items, and partly to the decision-maker’s mental state. We expected the same pattern of results if a decision-maker made a choice slowly. Finally, we expected that there would be an interaction effect between response time and carefulness/threshold. Specifically, we predicted that participants would learn more about a decision-maker’s preferences if they were cautious (where a fast decision versus a slow decision would be more meaningful) compared to if they were careless (where a fast decision versus a slow decision would be less meaningful or indicative of their preferences).

Figure 7. Schematic describing the DDM’s threshold parameter, and the inference task in Experiment 3. (A) The DDM’s parameter “threshold” (θ) describes how careful a decision-maker is. If a decision-maker has a fixed drift rate (meaning that their perceived value difference or preference for one item over the other is the same), a more cautious decision-maker (having a high threshold) will take longer to make a choice than a less cautious decision-maker (having a low threshold). (B) In an inverted DDM, if the observer does not know the decision-maker’s drift rate (preferences) or threshold (carefulness), the observer must infer whether a decision was fast because the decision-maker had a strong preference for that item, or because the decision-maker was being careless and not thinking too much about their decision. In Experiment 3, participants are told the decision-maker’s carefulness and response time and must infer their strength of preference.

Methods

Participants. 481 participants with U.S. IP addresses were recruited through Amazon Mechanical Turk. Participants were paid \$1.00 in compensation. In our

preregistration, we indicated the following exclusion criteria: Participants were excluded from the study if they answered more than two out of eight manipulation check questions incorrectly. (These manipulation check questions directly queried whether the participants had read the stimuli text. Answering them incorrectly demonstrated lack of attention.) We determined the sample size via a power analysis based on the effect size of pilot studies and these exclusion criteria. In this power analysis, we calculated that we would need at least 27 participants to achieve 99% power with our pilot result effect size and $\alpha=0.01$, using a Type-III F -test, but to ensure clear results chose to use a 480 as our sample size (one extra participant was recruited by the platform). However, our original exclusion criteria resulted in 212 participants being excluded (44% of our participants), which felt overly restrictive. For that reason, we decided to not use any exclusion criteria in Experiment 3, thereby including all participants in the analysis (we discuss the manipulation checks further in the Discussion). When we redid the power analysis calculation based on the pilot data with no exclusion criteria, using the same parameters as above, we calculated that we would need at least 179 participants to achieve 99% power. Thus, since 481 participants were included in the analysis, we still included a sample size with the desired power.

In the Supplemental Material, we restrict our sample to that specified by the preregistered exclusion criteria, and present the planned analyses and figures based on this sample for comparison. The results were similar to those presented here (Figure S2).

Stimuli and Procedure. Video stimuli were the same as in Experiment 1.

After consenting, participants read introductory text: “In this experiment, you will see videos of someone choosing between an item in one bowl and an item in another bowl. There are 8 videos. Please do your best to answer the questions afterward.” Participants then entered the main experiment.

Participants either read descriptive text emphasizing that the decision-maker was feeling meticulous and cautious (high threshold condition), or that the decision-maker was feeling careless and inattentive (low threshold condition). The text for “high threshold” read:

“This person is choosing between two items. It is very important to them to be very careful about the following decision. This person cares a lot about being *meticulous and cautious* right now. They had an easy day at work and so do not have much on their mind besides this decision.” The text for “low threshold” read: “This person is choosing between two items. It is not important to them to be careful about the following decision. This person is fine being *careless and inattentive* right now. They had a hard day at work and have a lot on their mind besides this decision.”⁵

Participants then saw a video underneath this text showing a decision-maker making a decision between two items, A and B, where the decision-maker took 3 seconds from video onset or 9 seconds from video onset to make these decisions. Next, participants were asked “How does the person feel right now?” and were given a binary choice between a “Meticulous and cautious” radio button and a “Careless and inattentive” radio button, which was used as a manipulation check to ensure the participants had read the descriptive text (the ordering of these buttons was random on each trial). Participants were finally shown the critical question, asking about the decision-maker’s preferences: “How much do you think they value their chosen item?” Participants moved a slider containing values from 0 (labeled “Neutral between items”) to 100 (labeled “Strongly prefers item”; 50 was labeled “Moderately prefers item”). No grid lines were shown, nor the numbers 0-100. The slider was initially set at 0, and participants were required to click and move the slider, but they could move it back to 0 if desired. After answering these two questions, participants could advance to the next page to see the next set of text/video/questions (Figure 8). There were eight pages total. All of the text/video/questions were on the same page, and participants could replay videos as often as they wished. Participants could not return to a previous page.

⁵ This text was designed to isolate the threshold parameter as much as possible, and avoided e.g. mentioning differences in the absolute value of items, or response time cutoffs or pressure, for this reason. The information about having a hard day at work was included because the decision-maker appeared very focused in all video stimuli, and describing them as “careless” appeared unnatural otherwise.

This person is choosing between two items. It is very important to them to be very careful about the following decision. This person cares a lot about being meticulous and cautious right now. They had an easy day at work and so do not have much on their mind besides this decision.

How does the person feel right now?

Meticulous and cautious

Careless and inattentive

How much do you think they value their chosen item?

Neutral between items

Moderately prefers item

Strongly prefers item

>>

Figure 8. Experiment 3 Stimulus. Participants watched a 3- or 9-second video of the decision-maker choosing between A and B then answered questions.

Each participant was randomly assigned to one of 12 counter-balancing variants in which the labeling and orientation of viewed videos were held fixed, and within which the participant viewed eight videos varying in response time and threshold. These eight videos presented in random order: Half of the choices were made in 3 seconds and half in 9 seconds (response time counterbalancing), half with the “high threshold” text and half with the “low threshold text” (threshold counterbalancing), and half with the decision-maker’s right hand

and half with the left hand (choice counterbalancing). After watching all eight videos and answering the accompanying questions, participants had the opportunity to write comments then exited the survey.

Model. We assume, as in Experiment 1, that participant responses are based on the posterior mean estimate of the difference in utilities of the two items. We additionally assume that the carefulness manipulation influences participants’ assumed threshold, θ , such that those in the cautious condition will make predictions consistent with a high θ and those in the careless condition will make predictions consistent with a low θ . To capture this, we held fixed the β and α parameters fit in the previous two experiments, and fit θ separately to the two groups. We accounted for the different scale in the response slider (specific to Experiment 3) by setting the predicted response to $2 * \alpha \mathbb{E}[u_a - u_b \mid a \succ b, t]$. Besides these differences, the predictions were produced in the same way as for Experiment 1, using Equation 5.

Results

In Experiment 3, we investigated whether participants could make inferences about decision-makers’ preferences based on both timing and threshold (how careful the decision-maker is) simultaneously. Our model predicts a main effect of response time (the shorter the response time, the larger the inferred value difference between the two items will be) and a main effect of threshold (the higher the threshold, the larger the inferred value difference between the two items will be). Our model also predicts an interaction effect between response time and threshold, such that decision-makers who are cautious and fast to make decisions are predicted to particularly value their chosen item, compared to decision-makers who are cautious and slow to make decisions (whereas we expect a smaller difference for the response time variable if decision-makers are emphasized to be careless). Results are shown in Figure 9.

To analyze the results, we constructed a linear mixed effects model, with response time

Figure 9. Experiment 3 Results. (a) Experimental Results. The mean $\pm$ SE of the inferred value difference, $n=481$, is shown for the high threshold (“cautious,” blue) and the low threshold (“careless,” black) conditions, for each of the 3 and 9 second response time conditions. Responses are connected by threshold condition for emphasis. Individual participants’ averaged values for each threshold/time pair are also shown and connected. 0 represents the participant feeling that the decision-maker was neutral between items, 50 that the decision-maker moderately preferred their chosen item, and 100 that the decision-maker strongly preferred their chosen item. (b) Model Predictions.

and threshold as categorical fixed effects, subject as a random effect, and participants’ judgments of value (0-100, where 0 is “Neutral between items,” 50 is “Moderately prefers item,” and 100 is “Strongly prefers item”) as the continuous dependent variable. The estimate (beta parameter) for the main effect of response time was 4.07 (95%CI=[2.62,5.53], $t(3364.0)=5.48$, $p=4.5e-8$, t -tests using Satterthwaite’s method, $\eta_p^2=0.009$). The estimate for the main effect of threshold was 11.50 (95%CI=[10.04,12.95], $t(3364.0)=15.48$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.07$). The estimate for the interaction effect of response time and threshold was 2.62 (95%CI=[-0.29,5.53], $t(3364.0)=1.76$, $p=0.08$, t -tests using Satterthwaite’s method, $\eta_p^2=0.0009$): a marginal but not significant effect. The

estimated power of the interaction effect for $\alpha=0.01$ was 19.50% (95%CI=[14.25,25.68]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R package “simr.” (The estimate for the fixed effect intercept was 61.13, 95%CI=[59.69,62.59], $t(480.0)=82.97$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.94$.) Results were similar, using the preregistered exclusion criteria; these results are included in the Supplemental Material (Figure S2). Thus, while our empirical results matched our model predictions in that we observed the expected main effects of response time and threshold, we failed to observe the anticipated interaction effect.

Discussion

In Experiment 3, we asked whether participants could infer a decision-maker’s preferences from their choices and response times while incorporating information about their mental state of being cautious or careless. As expected, we observed a main effect of response time in the expected direction (spending less time on a decision was associated with larger inferred preferences), though the main effect of response time was weaker than in Experiment 1, when there was only one manipulation (response time) rather than two (threshold and response time) and the critical question had a different format. We also observed a main effect of threshold in the expected direction (more carefulness was associated with larger inferred preferences). Finally, our model predicted an interaction effect, wherein a greater value difference was expected for cautious decision-makers with different response times compared to careless decision-makers, which we did not observe.

To manipulate threshold, we included a textual description above our video stimuli describing whether the decision-maker was feeling cautious or careless. However, in the videos, the decision-maker always appeared to be focused and therefore cautious, and participants expressed confusion in the experiment’s comments section about the discrepancy between the textual description of the decision-maker’s mental state and the decision-maker’s facial expressions. This issue was reflected in the manipulation check

questions, which asked “How does the person feel right now?” with the binary options of “Meticulous and cautious” and “Careless and inattentive.” Participants asymmetrically answered the manipulation check questions incorrectly. When the described mental state was “Meticulous and cautious,” 387/481 (80%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 3.3/4). When the described mental state was “Careless and inattentive,” 261/481 (54%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 2.6/4).

In our preregistration, we had planned to exclude participants who answered more than two of the eight manipulation checks incorrectly, but these criteria would have resulted in 212 participants (44%) being excluded.⁶ We thus included all participants in the experiment, though results were similar when the exclusion criteria were applied (Figure S2). We expected that the high failure rate for the manipulation check exclusion criteria was due to the mismatch between the textual description of the decision-maker’s mental state and the facial expression confound in the associated videos, so conducted Experiments 4A and 4B to address this confound.

Experiments 4A and 4B: Sensitivity to a decision-maker’s mental state with different stimuli

In exploring whether knowing a decision-maker’s mental state of cautious or careless influences people’s inferences of their preferences, Experiment 3 introduced a confound in the threshold manipulation whereby the decision-maker’s facial expressions did not match the textual descriptions of their mental state. We sought to address this using different stimuli in Experiment 4A (text message videos) and Experiment 4B (vignettes). In Experiment 4A, participants saw text message videos between the decision-maker and a friend, Alice. Alice

⁶ Note the locations of the manipulation check options were random for each trial, which likely contributed to some attention-based mistakes across both threshold conditions, separate from the described facial expression confound.

either described the decision-maker as having a good day at work and careful in their decision-making (high threshold condition), or having a hard day at work (low threshold condition). Alice then asked the decision-maker to choose between items A and B. In the fast response condition, the decision-maker responded “Hm,” then generated a decision. In the slow response condition, the decision-maker responded “Hm,” then spent time deliberating, designated by a moving ellipsis, before giving a decision. In Experiment 4B, rather than watching a video, participants read a vignette of the decision-maker’s mental state, the time they took to make a choice, and their choice. The manipulation check questions in Experiments 4A and 4B remained the same as those in Experiment 3. In both experiments, we addressed the same experimental question as in Experiment 3; thus, the model predictions are qualitatively the same.

Methods

Participants. In Experiment 4A, 480 participants were recruited through Prolific and paid \$1.25 in compensation. In Experiment 4B, 481 participants were recruited through Prolific and were paid \$1.00 in compensation. We did not exclude any participants. Experiments 4A and 4B used the same sample size and analyses as Experiment 3, and were not preregistered. Both experiments were completed at a later date and on a different platform (Prolific instead of Amazon Mechanical Turk), so it is possible but unlikely that participants from the first round of experiments (pilots, Experiments 1-3) participated in the second round of experiments (pilots, Experiments 4A and 4B). Participants within the second round of experiments could only complete one of that round’s pilot experiments or main experiments.

Experiments 4A and 4B were IRB-approved by Princeton University, Protocol ID: 10859, Protocol Title: Computational Cognitive Science. Participants gave informed consent.

Stimuli and Procedure. The stimuli and procedure were similar to those in Experiment 3, but the main stimuli and critical question were changed. The main stimuli in

Experiment 4A were changed to text message videos, and in Experiment 4B to vignettes. The critical question for Experiments 4A and 4B was posed as in Experiment 1 (“What are the person’s feelings about item A and item B?”) rather than Experiment 3 (“How much do you think they value their chosen item?”), since we considered the output of a scale spanning “Strongly Prefers A,” “Neutral Between A and B,” and “Strongly Prefers B” to be more robust than Experiment 3’s scale spanning “Neutral between items,” “Moderately prefers [chosen] item,” and “Strongly prefers [chosen] item.”

After consenting, participants read introductory text. In Experiment 4A: “In this experiment, you will see text message videos of someone choosing between two items. There are 8 videos. Please do your best to answer the questions afterward.” In Experiment 4B: “In this experiment, you will read vignettes of someone choosing between two items. There are 8 vignettes. Please do your best to answer the questions afterward.” Participants then performed eight trials of the experimental task.

In Experiment 4A, on each trial participants read: “This person is choosing between two items. Please watch the video.” A video underneath showed a text message conversation between Alice and the decision-maker. In the low threshold condition, the text in the video was the following. Alice: “I know you had a rough day, sorry to ask you this when you’re distracted” / “but” / “of the two items we discussed” / “did you want” / “A or B?” The decision-maker then replied: “Hm” / “A” (or “B”). The slashes here represent line breaks. In the high threshold condition, the first two lines from Alice were replaced with: “I heard you had a nice day at work, so figured I’d ask you today since I know you’re careful about this stuff :)”. Pauses were placed before each new line to approximate typing and reading speed, and the total time of the videos was the same for the low and high threshold conditions. The response time manipulation was in how long the decision-maker took to answer “A” or “B.” In the fast response condition, the answer came 2.0 seconds after “Hm” appeared. In the slow response condition, the answer came after 8.9 seconds; during this extended pause, an ellipsis (...) periodically displayed. An example stimulus is shown in Figure 10.

Figure 10. Experiment 4A Stimulus. Participants watched a text message video of a decision-maker choosing between A and B quickly or slowly, then answered questions. This image shows the final frame of the video.

In Experiment 4B, participants read: “Your friend had an easy day at work and enjoys being meticulous and cautious about their decisions. You’ve just asked them which of two items they’d prefer: A or B. After awhile, they answer A.” This text describes the “high threshold” condition, and the slow response condition. In the “low threshold” condition, the first sentence was replaced with “Your friend had a hard day at work and is feeling careless and inattentive about their decisions.” In the fast response condition, the last sentence was replaced with “Without pausing, they immediately answer A.” Participants

saw a total of eight vignettes: high/low threshold $\times$ slow/fast response $\times$ A/B choice. An example stimulus is shown in Figure 11.

Your friend had a hard day at work and is feeling careless and inattentive about their decisions. You've just asked them which of two items they'd prefer: A or B. After awhile, they answer A.

How does the person feel right now?

Meticulous and cautious ☐ Careless and inattentive ☐

What are the person's feelings about item A and item B?

Strongly Prefers A Neutral Between A and B Strongly Prefers B

>>

Figure 11. Experiment 4B Stimulus. Participants read vignettes of the decision-maker choosing between A and B quickly or slowly, then answered questions.

Participants were then asked the manipulation check question: “How does the person feel right now?” Participants were given a binary choice between a “Meticulous and cautious” radio button and a “Careless and inattentive” radio button; the ordering of these buttons was randomized on each trial. Participants were finally shown the critical question, asking about the decision-maker’s preferences: “What are the person’s feelings about item A and item B?” Participants moved a slider containing values from 0 (labeled “Strongly Prefers A”) to 100 (labeled “Strongly Prefers B”; 50 was labeled “Neutral Between A and B”). No grid lines were shown, nor the numbers 0-100. The slider was initially set at 50, and participants were required to click and move the slider, but they could move it back to 50 if desired. After answering these two questions, participants could advance to the next page. There were eight pages total. All of the text/(video)/questions were on the same page, and participants could replay videos as often as they wished if videos were present. Participants could not return to a previous page.

Each participant was randomly assigned to one of six counter-balancing variants in

which the item labels were held fixed (AB, BA, CA, AC, BC, and CB). Only the AB variant, in which the items were labeled A and B, will be described for expository purposes. Within each variant the participant viewed eight text message videos (Experiment 4A) or eight vignettes (Experiment 4B) varying in response time and threshold. These eight trials were presented in random order. Participants answered two questions for each of the combinations [cautious/fast], [careless/fast], [cautious/slow], [careless/slow]. After watching all eight text message videos or vignettes and answering the accompanying questions, participants had the opportunity to write comments then exited the survey.

Results

Here we asked whether people would incorporate a decision-maker’s carefulness in inferring their preferences from response times, using different stimuli from Experiment 3: text message videos in Experiment 4A, and vignettes in Experiment 4B. The model predictions are identical to those for Experiment 3: we predicted a main effect of response time (the shorter the response time, the larger the inferred value difference between the two items) and a main effect of threshold (the higher the threshold, the larger the inferred value difference between the two items). We also predicted an interaction effect between response time and threshold, such that decision-makers who were cautious and fast to make decisions would be judged to particularly value their chosen item, compared to decision-makers who were cautious and slow to make decisions (whereas we would expect a smaller difference for the response time variable if decision-makers were emphasized to be careless). Results are shown in Figure 12.

To analyze the results, we created a linear mixed effects model, with response time and threshold as categorical fixed effects, subject as a random effect, and participants’ preference judgments (0-100, 50 was “Neutral Between A and B,” 0 was “Strongly Prefers A” and 100 was “Strongly Prefers B”) as the continuous output variable.

For Experiment 4A, the results from the linear mixed effects model were as follows.

Figure 12. Experiment 4 Results. (a) Experiment 4A (Text Message Videos), $n=480$, and (b) Experiment 4B (Vignettes), $n=481$. The mean $\pm$ SE of inferred value difference is shown for the high threshold (“cautious,” blue) and the low threshold (“careless,” black) conditions, for each of the fast and slow response time conditions. Responses were reoriented to be 0-50: 0 indicates the participant felt the decision-maker was neutral between items A and B, and 50 indicates the participant felt the decision-maker strongly preferred the chosen item (A or B, whichever was appropriate to the trial). Responses are connected by threshold condition for emphasis. Individual participants’ averaged values for each threshold/time pair are also shown and connected.

The estimate (beta parameter) for the main effect of response time was 10.49 (95%CI=[9.74,11.24], $t(3357.0)=27.36$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.18$). The estimate for the main effect of threshold was 2.50 (95%CI=[1.75,3.26], $t(3357.0)=6.53$, $p=8e-11$, t -tests using Satterthwaite’s method, $\eta_p^2=0.01$). The estimate for the interaction effect of response time and threshold was 1.59 (95%CI=[0.086,3.09], $t(3357.0)=2.07$, $p=0.038$, t -tests using Satterthwaite’s method, $\eta_p^2=0.001$): a significant effect. The estimated power of the interaction effect for $\alpha=0.01$ was 30.00% (95%CI=[23.74,36.86]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R

package “simr.” (The estimate for the fixed effect intercept was 23.78, 95%CI=[22.84,24.73], $t(479.0)=49.58$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.84$.) Thus, our empirical results matched our model predictions in that we observed the expected main effects of response time and threshold. We also observed the predicted interaction effect (though only at 30.00% estimated power).

Participants were also asked a manipulation check question: “How does the person feel right now?” Participants had to choose between “Meticulous and cautious” and “Careless and inattentive.” (The locations of these options was random on every trial, which likely contributed to some mistakes.) However, participants asymmetrically answered the manipulation check questions incorrectly in Experiment 4A. When the described mental state was “Meticulous and cautious,” 321/480 (67%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 2.9/4). When the described mental state was “Careless and inattentive,” 121/480 (25%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 1.8/4). Under the preregistered exclusion criteria specified for Experiment 3, 341 participants (71%) would have been excluded for answering more than two of the eight manipulation checks incorrectly. Thus, we elected to not exclude participants based on their performance on the manipulation check (nor for any other reason).

For Experiment 4B, the results from the linear mixed effects model were as follows. The estimate (beta parameter) for the main effect of response time was 4.81 (95%CI=[3.71,5.91], $t(3364.0)=8.54$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.02$). The estimate for the main effect of threshold was 9.03 (95%CI=[7.93,10.14], $t(3364.0)=16.06$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.07$). The estimate for the interaction effect of response time and threshold was 0.79 (95%CI=[-1.42,2.99], $t(3364.0)=0.70$, $p=0.48$, t -tests using Satterthwaite’s method, $\eta_p^2=0.0001$): this effect was not present. The estimated power of the interaction effect for $\alpha=0.01$ was 2.50% (95%CI=[0.82, 5.74]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R package “simr.”

(The estimate for the fixed effect intercept was 27.66, 95%CI=[26.67,28.64], $t(480.0)=54.94$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.86$.) Thus, while we observed the expected main effects of response time and threshold, the predicted interaction effect was not present in Experiment 4B.

Participants did not answer the manipulation check questions asymmetrically in Experiment 4B. When the described mental state was “Meticulous and cautious,” 369/481 (77%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 3.38/4). When the described mental state was “Careless and inattentive,” 379/481 (79%) of participants answered at least 3 out of the 4 trials correctly (the mean correct number of trials was 3.44/4). Under the preregistered exclusion criteria specified for Experiment 3, 101 participants (21%) would have been excluded for answering more than two of the eight manipulation checks incorrectly. We did not exclude these participants, however, for consistency with Experiment 4A.

Discussion

The pattern of results across Experiments 3, 4A, and 4B support our model’s prediction that people can infer decision-makers’ preferences based on response times in a way that is appropriately sensitive to the decision-maker’s mental state. First, we observed a main effect of response time across all three experiments, even though the response time manipulation was implemented very differently for each (time for a decision-maker to visually reach for a bowl in Experiment 3, time to type a text message response in Experiment 4A, and “After awhile, they answer A,” or “Without pausing, they immediately answer A,” in Experiment 4B). The main effect of response time was strongest in Experiment 4A, followed by Experiment 4B, and weakest in Experiment 3. One difference between Experiment 3 and Experiments 4A and 4B is that Experiment 3 had a different critical question, but even if it had had the same critical question, Experiments 4A and 4B had distinct enough patterns of results from each other (response time, threshold, interaction

effects) that it is hard to know how this switch would have affected the results. Regardless, we observed a significant effect of response time across Experiments 1, 2, 3, 4A, and 4B, whereby participants inferred that a decision-maker had a stronger preference for an item the less time the decision-maker spent choosing it, supporting the robustness of the effect.

We also observed a main effect of threshold across Experiments 3, 4A, and 4B, in that if the decision-maker was described as cautious, participants inferred that the decision-maker had a stronger preference for their chosen item. This threshold effect was less present in Experiment 4A compared to Experiments 4B and 3, likely because the threshold manipulation was less explicitly emphasized (participants had to infer the decision-maker’s mental state from Alice’s description). However, we find it reassuring that a threshold manipulation as minor (and naturalistic) as that in Experiment 4A was enough to influence people’s inferences about the decision-maker’s preferences. We believe that the presence of both main effects—threshold and response time—across quite different paradigms speaks to the strength of people’s ability to socially infer preferences from response times while taking into account decision-makers’ mental states.

The model also predicted an interaction effect, wherein the effect of response time on inferred value is larger for cautious vs. careless decision-makers. This interaction effect did not occur at all in Experiment 4B ($p=0.48$, 2.5% power), was only marginally significant in Experiment 3 ($p=0.08$, 19.5% power), and was significant at low power in Experiment 4A ($p=0.04$, 30% power). Further work will be required to uncover how people jointly reason about how careful a decision-maker is and their response times in inferring the decision-maker’s preferences.

Finally, manipulation check performance was strongest and symmetric across threshold conditions in Experiment 4B, worse and asymmetric in Experiment 3, and poorest and asymmetric in Experiment 4A. Experiments 4B and 3 contained explicit threshold manipulations so it is not surprising that their manipulation check performance was higher than in Experiment 4A, in which the threshold manipulation was weaker. It was also

expected that Experiment 4B would have higher and more symmetric manipulation check performance compared to Experiment 3, because Experiment 4B did not contain a facial expression confound. It is surprising, however, that Experiment 4A had asymmetric performance across the threshold conditions because the stimuli were texts without a facial expression confound. One explanation is that the high threshold condition in Experiment 4A contained the word “careful” (which is semantically close to the high-threshold manipulation check option “Meticulous and cautious”), while the low threshold condition only included the word “distracted” (which is arguably less close to the low-threshold manipulation check option “Careless and inattentive”). Alternatively, perhaps our stimuli generally resulted in a bias for attributing carefulness to the decision-maker; Convincing participants that the decision-maker was being at least somewhat deliberate in their choices was a requisite property in our experiments. To conclude, the fact that we observed main effects of response time and threshold across all three experiments suggests that manipulation check performance is informative but not central to our model’s predictions and results.

General Discussion

If your friend takes a long time to decide between two options, you can infer that both options have similar value to her. Making these kinds of inferences about others’ preferences is ubiquitous in daily life as people seamlessly integrate choices and response times. In this work, we created a computational model to describe this phenomenon by inverting a classic generative model of decision-making, the Drift Diffusion Model. We tested the predictions of this model for three experimental questions.

In Experiment 1, participants inferred a decision-maker preferred a chosen item more when they spent a longer time deciding between two items, matching model predictions. In Experiment 2, participants inferred a decision-maker’s relative preferences from two choices and response times to infer a decision-maker’s preference about a third unseen choice, and they inferred the model’s predicted choice greater than chance across all conditions.

Experiment 2 also included two control conditions in which participants only needed choice information to anticipate the decision-maker’s unseen choice. As predicted, participants thought the predicted choice was more likely in these control conditions, providing an upper bound for participant likelihoods in this experiment. On the other hand, participant behavior for the different conditions was more variable than expected, which should be incorporated into the model in future work. In Experiment 3, we asked whether participants could integrate the decision-maker’s mental state of cautiousness or carelessness when inferring preferences. We observed the expected main effects of threshold and response time in three experiments spanning a range of stimuli (Experiments 3, 4A, and 4B), but a more complicated pattern of results in observing the expected interaction effect. Future work should investigate how experimental manipulations of threshold and response times are entangled in inferences about decision-makers’ preferences.

The uniqueness of this work lies in the question asked and the inversion of the DDM to formally characterize our predictions. Specifically, we asked how people infer other’s preferences by observing their response times and choices, which involves characterizing people’s theory of mind surrounding response time. The work closest to this appears in a recent preprint by Konovalov and Krajbich (2020), who demonstrate that people use response times to make inferences about others’ preferences, and further than people apply this knowledge strategically in bargaining games. Where our work differs is we formally describe these inferences about other’s minds by inverting the DDM and making direct, quantitative predictions, whereas Konovalov and Krajbich (2020) make qualitative predictions based on the DDM. (Konovalov and Krajbich (2020) also apply their work to strategic settings, while we explore prediction for novel choices and integrating the decision-maker’s mental state of cautiousness or carelessness.)

Our work has several potential extensions. Previous work already touches on interesting applications in strategic gameplay and mechanism design: Frydman and Krajbich (2018) designed an experiment in which some participants were explicitly shown others’

response times for decision-making, and showed that those participants used their response time inferences about others’ preferences to achieve better task outcomes compared to participants who did not have access to response time information. Expanding upon this, Konovalov and Krajbich (2020) showed that participants improved their outcomes by integrating either explicitly-presented or real-time perception of response time in a bargaining game. There is a fascinating aspect of mechanism design in these studies, in that people’s performance is improved by the researchers introducing additional information—response times—from which players make inferences. This aspect of mechanism design is even more prominent in Krajbich et al. (2014),⁷ albeit Krajbich et al. (2014) did not involve participants inferring other people’s preferences. Krajbich et al. (2014) expected based on DDM predictions that participants would spend more time on difficult choices—in which options were of similar value—even when opportunity costs suggested participants should pick one and move on. The authors designed an experiment where opportunity cost was high, and then implemented an intervention to end trials early when participants took too long to respond (an intervention they anticipated would target difficult choices in which participants were close to indifference). As predicted, after experiencing this intervention participants went on to earn higher rewards on the task on non-intervention blocks, from which the authors conclude that “it may be possible to improve people’s welfare with simple interventions or behavioral training.” In our work, we created an inverse DDM that generates predictions about how people infer other’s preferences from response times. One could imagine that this output could be a valuable tool in mechanism design: Presenting the output from such a model in strategic games, either as a training intervention or throughout the whole task, could serve as an assistive tool to improve people’s performance. Integrating mechanism design for inference of others’ preferences from response times in real world settings would be an intriguing application for this work.

The inverted DDM captures some elements of social inference from response times, but

⁷ See also Oud et al. (2016).

a useful future direction would be to incorporate the idea of the overall value of choice sets, and other elements of “context.” For example, the DDM describes a decision-maker’s difference in value between two items, but it does not incorporate the idea that choice sets can be of overall high value (e.g. buying a house, going to college) or low value (e.g. what candy bar to eat). Empirically, as the overall value of a choice set increases, response times decrease; fortunately, researchers have developed alternative sequential sampling models that can account for this effect (Clithero, 2018b; Shevlin & Krajbich, 2020). Incorporating “context” factors in inverted DDMs such as overall value, choice option similarity (Bhatia & Mullett, 2018), attention (Krajbich, 2019) and (as investigated in this work) mental states can not only help us describe and predict people’s behavior, but could be used as tools to help nudge people towards more optimal use of their time.

Finally, the particular inversion of the DDM in this work has applications for human-robot interaction. This work could be interpreted as inverse reinforcement learning (Jara-Ettinger, 2019; Ng & Russell, 2000) with algorithm run times: inferring someone’s utility function based on how long it takes them to decide. Artificial decision-makers are generally quite poor at social inference, especially compared to human performance, and so incorporating formal models how people learn about others through response times will only improve communication between people and automated systems.

Conclusion

If your friend immediately says “Yes!” to a movie but contemplates before saying “...yes” to hiking, you’ve learned something about the relative strength of her preferences that you couldn’t learn just from her choices. People constantly use response time to make inferences about each other’s preferences in daily life, and in this work we sought to quantitatively capture and predict that phenomenon. We present a rational model of inferring preferences from response time, using the Drift Diffusion Model to describe how preferences influence response time and Bayesian inference to invert this relationship. We

use our inverted DDM to predict participant behavior for three experimental questions. We first demonstrated that people could infer that a decision-maker preferred a chosen item more if they spent longer deliberating (Experiment 1). We then showed that people could predict a decision-maker’s choice in a novel comparison based on inferring the decision-maker’s relative preferences from previous response times and choices (Experiment 2). Finally we observed that people could incorporate information on a decision-maker’s mental state, cautiousness or carelessness, in inferring a decision-maker’s preferences from their response times (Experiments 3, 4A, and 4B).

We live in a social world, and people perform a startling number of unconscious inferences in navigating others’ preferences and goals. While economists and psychologists have been characterizing how we learn about others’ minds from their choices for decades, capturing the more subtle inferences made from watching someone think is also essential to understanding how we are so good at knowing what is not said. In this paper we advance this goal by creating a rational model of inverted decision-making to describe these inferences from response times. Much remains to be incorporated into these models. Within our results, more remains to be investigated around how preferences are inferred if decision-makers’ mental states are presented in ways that interact with cues like response time. Within DDMs, response times are influenced by other factors than just value difference, such as overall value of choice sets and many different types of “context” (we investigated one version of context, a decision-maker’s mental state of cautiousness or carelessness). Outside of DDMs, incorporating tone in addition to choices and response times will be key to describing how people make inferences about others’ preferences. Fortunately, the applications for formalizing models of social inferences are exciting and broad. People already draw impressive conclusions about others’ beliefs, preferences, and future actions from their choices and response times, and fully elucidating models of these social inferences not only helps us understand the mechanisms behind such powerful computational process, but may help us help people make better choices—via mechanism

design, or building tools based on humans' existing capacities—in the future.

Acknowledgements. This work was funded in part by John Templeton Foundation [grant number 61454] to T.L.G. Declarations of interest: none.

Author contributions. All authors developed and contributed to the study design, provided critical feedback throughout the project, and contributed to figure design and writing. Experimental data collection and analysis were performed by V.G. F.C. developed and analyzed the model, and drafted model-centric parts of the manuscript. V.G. drafted the rest of the manuscript with input from all authors. All authors approved the final version of the manuscript for submission.

Open Practices Statement. Experiments 1, 2 and 3 were preregistered (<https://osf.io/8n3kd>) and materials and data are available at <https://osf.io/pczb3>.

References

- Acuna, B. D., Sanes, J. N., & Donoghue, J. P. (2002). Cognitive mechanisms of transitive inference. *Experimental Brain Research*, 146(1), 1–10. <https://doi.org/10.1007/s00221-002-1092-y>
- Alós-Ferrer, C., & Garagnani, M. (2020). Strength of preference and decision making under risk. *University of Zurich, Department of Economics, Working Paper*, (330). <https://doi.org/10.2139/ssrn.3428515>
- Amasino, D. R., Sullivan, N. J., Kranton, R. E., & Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature Human Behaviour*, 3(4), 383–392. <https://doi.org/10.1038/s41562-019-0537-2>
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 1–10. <https://doi.org/10.1038/s41562-017-0064>
- Bhatia, S., & Mullett, T. L. (2018). Similarity and decision time in preferential choice. *Quarterly Journal of Experimental Psychology*, 71(6), 1276–1280. <https://doi.org/10.1177/1747021818763054>
- Bitzer, S., Park, H., Blankenburg, F., & Kiebel, S. J. (2014). Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Frontiers in Human Neuroscience*, 8, 102. <https://doi.org/10.3389/fnhum.2014.00102>
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700. <https://doi.org/10.1037/0033-295X.113.4.700>
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128), 95–98. <https://doi.org/10.1126/science.1231965>
- Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(3), 538. <https://doi.org/10.1037/0278-7393.11.3.538>

- Busemeyer, J. R., & Rieskamp, J. (2014). Psychological research and theories on preferential choice. *Handbook of choice modelling*. Edward Elgar Publishing.
<https://doi.org/10.4337/9781781003152.00008>
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432.
<https://doi.org/10.1037//0033-295X.100.3.432>
- Chabris, C. F., Laibson, D., Morris, C. L., Schuldt, J. P., & Taubinsky, D. (2008). *Measuring intertemporal preferences using response times* (tech. rep.). National Bureau of Economic Research. <https://doi.org/10.3386/w14353>
- Chabris, C. F., Morris, C. L., Taubinsky, D., Laibson, D., & Schuldt, J. P. (2009). The allocation of time in decision-making. *Journal of the European Economic Association*, 7(2-3), 628–637.
<https://doi.org/10.1162/jeea.2009.7.2-3.628>
- Clithero, J. A. (2018a). Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization*, 148, 344–375.
<https://doi.org/10.1016/j.jebo.2018.02.007>
- Clithero, J. A. (2018b). Response times in economics: Looking through the lens of sequential sampling models. *Journal of Economic Psychology*, 69, 61–86.
<https://doi.org/10.1016/j.joep.2018.09.008>
- Dai, J., & Busemeyer, J. R. (2014). A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *Journal of Experimental Psychology: General*, 143(4), 1489.
<https://doi.org/http://dx.doi.org/10.3758>
- Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology*, 41(3), 260–274.
<https://doi.org/10.1006/jmps.1997.1167>
- Diederich, A. (2003). Decision making under conflict: Decision time as a measure of conflict strength. *Psychonomic Bulletin & Review*, 10(1), 167–176. <https://doi.org/10.3758/BF03196481>
- Echenique, F., & Saito, K. (2017). Response time and utility. *Journal of Economic Behavior & Organization*, 139, 49–59. <https://doi.org/10.1016/j.jebo.2017.04.008>

- Frydman, C., & Krajbich, I. (2018). Using response times to infer others' beliefs: An application to information cascades. *Available at SSRN 2817026*. <https://doi.org/10.2139/ssrn.2817026>
- Genz, A. C., & Malik, A. A. (1980). Remarks on algorithm 006: An adaptive algorithm for numerical integration over an N-dimensional rectangular region. *Journal of Computational and Applied Mathematics*, 6(4), 295–302. [https://doi.org/10.1016/0771-050X\(80\)90039-X](https://doi.org/10.1016/0771-050X(80)90039-X)
- Gill, D., & Prowse, V. L. (2017). Using response times to measure strategic complexity and the value of thinking in games. *Available at SSRN 2902411*. <https://doi.org/10.1007/s10683-015-9463-y>
- Harris, M. R., & McGonigle, B. O. (1994). A model of transitive choice. *The Quarterly Journal of Experimental Psychology*, 47(3), 319–348. <https://doi.org/10.1080/14640749408401362>
- Hu, J., Lucas, C. G., Griffiths, T. L., & Xu, F. (2015). Preschoolers' understanding of graded preferences. *Cognitive Development*, 36, 93–102. <https://doi.org/10.1016/j.cogdev.2015.09.012>
- Jara-Ettinger, J. (2019). Theory of mind as inverse reinforcement learning. *Current Opinion in Behavioral Sciences*, 29, 105–110. <https://doi.org/10.1016/j.cobeha.2019.04.010>
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20(8), 589–604. <https://doi.org/10.1016/j.tics.2016.05.011>
- Jern, A., Lucas, C. G., & Kemp, C. (2017). People learn other people's preferences through inverse decision-making. *Cognition*, 168, 46–64. <https://doi.org/10.1016/j.cognition.2017.06.017>
- Kononov, A., & Krajbich, I. (2019). Revealed strength of preference: Inference from response times. *Judgment & Decision Making*, 14(4).
- Kononov, A., & Krajbich, I. (2020). Decision times reveal private information in strategic settings: Evidence from bargaining experiments. *Available at SSRN 3023640*. <https://doi.org/10.2139/ssrn.3023640>
- Krajbich, I. (2019). Accounting for attention in sequential sampling models of decision making. *Current Opinion in Psychology*, 29, 6–11. <https://doi.org/10.1016/j.copsyc.2018.10.008>

- Krajibich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292.
<https://doi.org/10.1038/nn.2635>
- Krajibich, I., Bartling, B., Hare, T., & Fehr, E. (2015). Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nature Communications*, *6*(1), 1–9.
<https://doi.org/10.1038/ncomms8455>
- Krajibich, I., Hare, T., Bartling, B., Morishima, Y., & Fehr, E. (2015). A common mechanism underlying food choice and social decisions. *PLoS Computational Biology*, *11*(10), e1004371.
<https://doi.org/10.1371/journal.pcbi.1004371>
- Krajibich, I., Oud, B., & Fehr, E. (2014). Benefits of neuroeconomic modeling: New policy interventions and predictors of preference. *American Economic Review*, *104*(5), 501–06.
<https://doi.org/10.1257/aer.104.5.501>
- Krajibich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857. <https://doi.org/10.1073/pnas.1101328108>
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, *4*, 863.
<https://doi.org/10.3389/fpsyg.2013.00863>
- Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., Markson, L., & Hu, J. (2014). The child as econometrician: A rational model of preference understanding in children. *PloS One*, *9*(3). <https://doi.org/10.1371/journal.pone.0092160>
- Maybery, M. T., Bain, J. D., & Halford, G. S. (1986). Information-processing demands of transitive inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(4), 600. <https://doi.org/10.1037/0278-7393.12.4.600>
- Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making*, *5*(6), 437.
<https://doi.org/10.2139/ssrn.1901533>

- Moffatt, P. G. (2005). Stochastic choice and the allocation of cognitive effort. *Experimental Economics*, 8(4), 369–388. <https://doi.org/10.1007/s10683-005-5375-6>
- Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In P. Langley (Ed.), *Proceedings of the 17th International Conference on Machine Learning* (pp. 663–670).
- Oud, B., Krajbich, I., Miller, K., Cheong, J. H., Botvinick, M., & Fehr, E. (2016). Irrational time allocation in decision-making. *Proceedings of the Royal Society B: Biological Sciences*, 283(1822), 20151439. <https://doi.org/10.1098/rspb.2015.1439>
- Polania, R., Krajbich, I., Grueschow, M., & Ruff, C. C. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron*, 82(3), 709–720. <https://doi.org/10.1016/j.neuron.2014.03.014>
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427–430. <https://doi.org/10.1038/nature11467>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59. <https://doi.org/10.1037/h0035486>
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4), 260–281. <https://doi.org/10.1016/j.tics.2016.01.007>
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, 117(523), 1243–1259. <https://doi.org/10.1111/j.1468-0297.2007.02081.x>
- Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5), 927–939. <https://doi.org/10.1016/j.neuron.2016.04.036>
- Shevlin, B., & Krajbich, I. (2020). Attention as a source of variability in decision-making: Accounting for overall-value effects with diffusion models. *PsyArXiv*. <https://doi.org/10.31234/osf.io/rewtq>

- Spiliopoulos, L., & Ortmann, A. (2018). The BCD of response time analysis in experimental economics. *Experimental Economics*, 21(2), 383–433.
<https://doi.org/10.1007/s10683-017-9528-1>
- Thayer, E. S., & Collyer, C. E. (1978). The development of transitive inference: A review of recent approaches. *Psychological Bulletin*, 85(6), 1327. <https://doi.org/10.1037/0033-2909.85>
- Towal, R. B., Mormann, M., & Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, 110(40), E3858–E3867. <https://doi.org/10.1073/pnas.1304429110>
- Wilcox, N. T. (1993). Lottery choice: Incentives, complexity and decision time. *The Economic Journal*, 103(421), 1397–1417. <https://doi.org/10.2307/2234473>
- Zhao, W. J., Diederich, A., Trueblood, J. S., & Bhatia, S. (2019). Automatic biases in intertemporal choice. *Psychonomic Bulletin & Review*, 26(2), 661–668.
<https://doi.org/doi.org/10.3758/s13423-019-01579-9>

Supplemental Material

Experiments 1, 2, and 3 were preregistered (<https://osf.io/8n3kd>) and materials and data are available at <https://osf.io/pczb3>.

Experiment 2, Additional Analysis

Figure S1. Experiment 2 Results, Additional Analysis. (a) Main Figure (see Figure 6). (b) Individual (Connected) Participant Responses. Note the prevalence of responses following a “50 - 50 - 100 - 100” pattern for the conditions “NeitherChosen - BothChosen - FastChoice - SlowChoice.”

Future work may also explore some quirks of the empirical results. In Figure 6, the results indicate that while participants did not often predict the non-predicted choice, very many participants chose 50 for the timing-inference based conditions, and 100 for the control conditions. (These numbers indicate participants’ inferred likelihood for the predicted choice; the critical question participants were asked was: “This person is now offered a choice between items B and C. What choice do you think they’d make, and how likely do you think it is that they’d make that choice?”) This pattern indicates that participants were not using response times when making their choices, and that choices alone could make them estimate 100% likelihoods for the decision-maker’s choice on an unseen choice. Figure S1 suggests that

a number of participants answered in this “50-50-100-100” sequence (for conditions “NeitherChosen-BothChosen-FastChoice-SlowChoice”), implying that if such responses had been discouraged the effect observed here may have been even stronger.

Experiment 3, Excluded Participants

The analysis was identical to Experiment 3 except that participants were excluded if they answered more than two out of eight manipulation check questions incorrectly, as was defined in the preregistration. Of 481 recruited participants, 212 met this exclusion criteria, so 269 participants were included. These results were similar to the Experiment 3 results with all participants included, and are described below (Figure S2).

From our linear mixed model, the estimate (beta parameter) for the main effect of response time was 6.58 (95%CI=[4.53,8.62], $t(1880.0)=6.31$, $p=4e-10$, t -tests using Satterthwaite’s method, $\eta_p^2=0.02$). The estimate for the main effect of threshold was 19.59 (95%CI=[17.54,21.63], $t(1880.0)=18.70$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.16$). The estimate for the interaction effect of response time and threshold was 3.65 (95%CI=[-0.43,7.74], $t(1880.0)=1.75$, $p=0.08$, t -tests using Satterthwaite’s method, $\eta_p^2=0.002$): not a significant effect. The estimated power of the interaction effect for $\alpha=0.01$ was 20.00% (95%CI=[14.69,26.22]) using a Type-III F -test from the R package “car,” generated via 200 simulations with the R package “simr.” (The estimate for the fixed effect intercept was 57.37, 95%CI=[55.67,59.08], $t(268.0)=66.15$, $p<2e-16$, t -tests using Satterthwaite’s method, $\eta_p^2=0.94$.)

Figure S2. Experiment 3 Results, with Excluded Participants. (a) Experimental Results.

The mean $\pm$ SE of the inferred value difference, $n=269$, is shown for the high threshold (“cautious,” blue) and the low threshold (“careless,” black) conditions, for each of the 3 and 9 second response time conditions. Responses are connected by threshold condition for emphasis. Individual participants’ averaged values for each threshold/time pair are also shown and connected. 0 represents the participant feeling that the decision-maker was neutral between items, 50 that the decision-maker moderately preferred their chosen item, and 100 that the decision-maker strongly preferred their chosen item. (b) Model Predictions.